2021

A COMPUTATIONAL APPROACH WITHIN MEDICAL RESEARCH

Ryan Christopher Hogan

A COMPUTATIONAL APPROACH WITHIN MEDICAL RESEARCH

A thesis submitted in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE

to the faculty of the

DIVISION OF COMPUTER SCIENCE, MATHEMATICS AND SCIENCE

of

THE LESLEY H. AND WILLIAM L. COLLINS
COLLEGE OF PROFESSIONAL STUDIES
at

ST. JOHN'S UNIVERSITY

New York

by

Ryan Hogan

Date Submitted ____________ Date Approved ___________

_________________________ _________________________

Ryan Hogan Dr. Christoforos Christoforou

ABSTRACT

A COMPUTATIONAL APPROACH WITHIN MEDICAL RESEARCH

Ryan Hogan

Within the context of medical image diagnosis, we explore novel computational models to facilitate the detection of two medical conditions that burden our society. In particular, this research focuses on the use of deep learning models for the detection of Alzheimer's Disease in Magnetic Resonance images (MRI) scans, as well as the detection of heart arrhythmias from electrocardiogram (ECG) recordings. We propose a novel architecture that depends on the 3D-CNN model to classify between MRI scans of cognitively healthy individuals and AD patients. Moreover, we explore the use of LSTM deep learning models to detect abnormal heart arrhythmias that present life-threatening challenges for individuals with underlying conditions that may not be recognized through current practices. The goal of this research is to measure the efficacy and predictability of applying deep learning techniques to detect AD by mapping the complex heterogeneity of the brain, and heart arrhythmias in ECG time-series recordings in a computational way.

ACKNOWLEDGMENTS

I wish to express my sincere appreciation for my advisor, Dr. Christoforos Christoforou, for his support and guidance throughout the culmination of this thesis. This work would not have been possible without his expertise in both computer science and neuroinformatic research.

Finally, I would like to thank my family, friends, and peers for all their love and encouragement throughout the pursuit of this study.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Deep Learning within Medical Research

The effort to introduce new forms of diagnosis and treatment options falls in the domain of an alternative approach to human intelligence. Deep learning techniques have offered a comprehensive understanding of the human body, as well as the identification of diseases and other health conditions alike. Current healthcare practitioners often ignore a computational approach that aid in the proper diagnosis of patients in a qualitative way [19]. By collecting informative data such as radiology scans, medical history information, and diagnostic data, deep learning offers a computational approach to derive probabilistic predictions by examining and correlating these interrelated concepts. In current practices, medical practitioners may examine health records and other patient information such as radiology scans and laboratory data to form an understanding of health conditions and provide diagnostic treatment options. However, it can be argued that the misdiagnosis of illnesses and health conditions tends to be detrimental to one's health due to inaccurate readings or human error, at an average rate of 10-15% of the time [10].

Due to such inaccuracies, deep learning provides a more sophisticated approach to assist doctors in the discovery and proper diagnosis of health conditions in a computational way. The collection of data from others who have experienced the same illnesses motivate a stronger prediction from deep learning models that achieve success by examining large amounts of information. Deep learning informs a probabilistic prediction through the collection of image and text-based qualitative data by learning “features” within a contiguous set of information [4]. In the case of an image, local neighborhoods within adjacent pixels can be examined to perform feature extraction in lower levels of the deep learning network. As more informative features are recognized throughout multiple iterations of the network, an output can be determined by adding non-linearity activation functions such as the sigmoid function. Medical research, with the help of deep learning, will continue to benefit from a computational approach of data collection and analysis to inform new insights into medical education and diagnosis.

1.2 Alternative Diagnosis to Alzheimer’s Disease and Arrhythmias

The most common method of tracking AD progression has been facilitated through magnetic resonance imaging (MRI) as a non-invasive tool for detecting changes in brain volumes throughout incremental scans. Machine learning has offered an alternate form of detection in MRI scans by utilizing multiple algorithms such as Support Machine Vector (SVM) and Random Forest (RF) to perform classification tasks through linear SVM classification and majority voting in RF. Other approaches such as Region Of Interest (ROI) patch detection, biomarkers, and cerebral spinal fluid (CSF) are alternative methods for performing classification of AD. However, there has been recent knowledge to suggest a potential lack of information in ROI patches as well as being error-prone and labor-intensive [16]. For this reason, ROI was not considered for this study despite being popular in some 3D models where whole-brain volumes are used, following a 3D patch for areas known to represent AD degeneration [16].

The detection of abnormalities within ECG data can often go unidentified because of the inability to receive proper healthcare or the subtleness of the arrhythmias. Because of this, a computational approach has been adopted to identify these potentially minute changes within time-series data to derive a diagnosis of arrhythmias within the heartbeat.

In this study, we propose the use of long short-term networks (LSTM) to provide an insightful approach in detecting irregularities in raw ECG data. LSTM neural networks, a subdivision recurrent neural networks (RNN), and deep learning do a remarkable job in remembering information from previous layers of the network while passing the most important information along to the next iteration of the network. In this case, LSTM is able to utilize the time-series data of ECG recordings to remember informative features in the signals and then pass the most relevant information along the network to detect changes between corresponding signals [15].

1.3 Thesis Statement

The goal of this research is to inform a reputable understanding of a computational approach to identify and diagnose health conditions in real-world applications. We propose two separate examples of how deep learning can facilitate new discoveries in medical research, particularly in the diagnosis of Alzheimer's Disease in MRI scans, and the detection of abnormal heartbeats in ECG recordings. Our question in this thesis is, "Does a deep learning approach offer discriminable insights to a proper diagnosis compared to current practices?". We hypothesize that the use of 3D convolutional neural networks to detect Alzheimer's Disease, and long short-term memory networks to identify abnormal heart arrhythmias will offer a more confident diagnosis to both conditions, including other health conditions in the domain through a computational approach.

CHAPTER 2

ALZHEIMER'S DETECTION THROUGH 3D CONVOLUTIONAL NEURAL NETWORKS

2.1 An introduction to Alzheimer's Disease

Alzheimer's disease (AD) is a debilitating type of dementia that can affect a person's memory and cognitive behavior in both mild and extreme cases such as the inability to remember, focus, or even understand who you are [26]. AD is a common disease that has no cure and worsens over time, typically occurring in individuals of 65 and older, making up 50-60% of all dementia cases [26]. The ability to cure AD is not currently known while treatments such as prescription medication and therapy have not shown to slow degeneration within the brain. Changes in the brain allow AD to damage neurons throughout multiple compartments dealing with memory and cognitive function, specifically in the limbic system, temporal lobe, hippocampus, and cerebral cortex. The damage of these neurons is not repairable and breaks the communication pathways that our brains use to create short-term memories and maintain continuous cognitive function.

The primordial development of AD can be seen through subtle changes in both behavioral and cognitive function, representing the transitional state of AD known as mild cognitive impairment (MCI). This is a critical stage of detection in which the disease can worsen and spread to other parts of the brain, slowing affecting regions that hinder a person's cognitive function depending on the degenerative location in the brain. Detection of AD in early stages (MCI) is an important task to develop treatment remedies and offer potential management to this overwhelming disease. The detection of MCI becomes a challenge due to the subtle physiological changes in brain volumes that can occur when a patient has not fully converted to AD. The predictive power in training may yield lower accuracies due to the difficulty required in identifying the conversion of MCI to AD. This can be seen in studies that do not meet similar benchmarks as the same model when classifying between NC and AD due to less discriminable brain changes. Model performance presented in [16] yielded an accuracy score of 73.04% when classifying between NC vs. AD and 68.49% in MCI vs. AD, respectively.

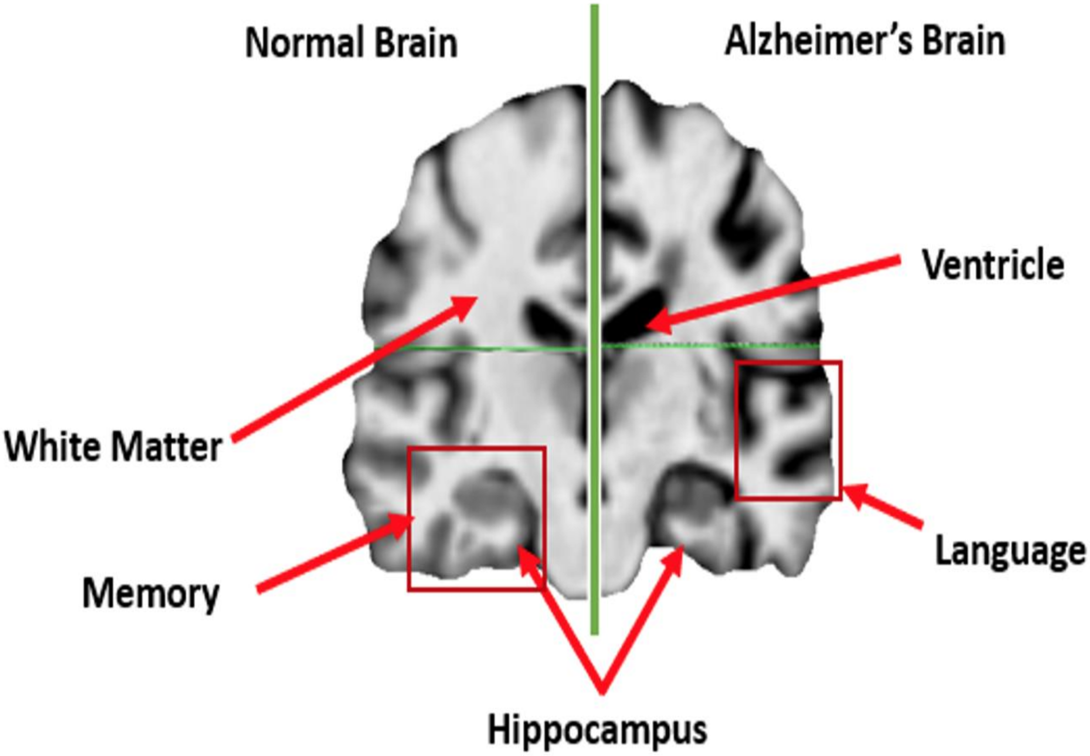

Figure 2.1: Left (healthy control) and Right (AD) brain comparison.

While under the same umbrella as machine learning, deep learning has initiated a more sophisticated approach by applying different methods such as convolutional neural networks (CNN), stacked auto-encoders (SAE), and deep belief networks (DBN) to extract low-level features within an image. For classification tasks, CNN's have proved to be a valuable application for extracting learnable parameters as well as feature mapping within medical imaging [14]. Moreover, the CNN explicitly takes an image as direct input to measure spatial information within adjacent pixels to assist in weight-sharing and backwards propagation. For this model, the CNN will be used to extract voxel information contained in MRI scans without the need for manual selection of image features typically done in ROI patches. However, to the best of our knowledge, most CNN approaches proposed in this domain, rely on 2D- convolutional layers that require MRI scans to be processed as 2D slices, thus ignoring the 3D spatial-structure of MRI brain scans. Processing MRI images as a 3D volume, instead of 2D slices, could potentially provide additional information that differentiates between AD, MCI, and HC groups, but it also introduces additional computational complexity in the model. In this paper, we explore the feasibility of a processing MRI image as a 3D volume, a CNN that utilized 3D convolutional layers.

The goal of this research can be split into two parts including: the efficacy of using a 3D CNN to detect discriminable changes within brain values between normal control and AD patients, as well as detecting changes between mild cognitively impaired and AD patients, respectively. The characteristics of a 3D convolutional network will allow for the use of all spatial dimensions when identifying changing within brain volumes, particularly in NC vs. AD patients. Our question in this thesis is, “Does a 3D convolutional neural network outperform 2D models that ignore the third spatial dimension?”. We hypothesize that 3D images have potential to detect degeneration between MRI scans of both HC and AD patients better than 2D slices by taking advantage of the third dimension, yielding a binary classification between two classes.

2.2 Related Work

Extensive research has been done to discovery of a cure to Alzheimer’s Disease, a disease that diminishes a total sense of sense due to its overwhelming symptoms and impairment on cognitive functions. The earliest studies of Alzheimer’s Disease have been established over 100 years ago, although our modern understanding of the disease has only been recognized since the early 1980’s [2]. In recent years, a new method of detection has been facilitated in a computational way to identify neurodegeneration in MRI scans. Researchers have begun to apply artificial learning techniques to bridge the gap in providing optimistic treatment options and encourage the discovery of a potential cure.

Early detection of Alzheimer’s Disease has been a common goal amongst researchers in the community. Identifying Alzheimer’s Disease in early stages may offer proactive treatment options before extensive degeneration occurs, causing atrophy in the brain. Leading research, such as [17], examines early detection of AD by combining a 2D convolutional neural network (CNN) with machine learning techniques such as Bagging [3] to identify discriminable brain regions that are known to represent Alzheimer’s Disease and mild cognitive impairment. By using the 2D CNN, an ensemble learning approach was employed to train base classifiers in 2D MRI slices. The ensemble learning approach allows the network to generalize the most discriminative MRI slices to predict classification results on unseen data. This study reports classification accuracies of 0.84 ± 0.05, 0.79 ± 0.04, and 0.62 ± 0.06, respectively, for classifying Alzheimer’s Disease vs Healthy Control (HC), mild cognitively impaired converters (MCIc) and HC, and MCIc vs MCInc (non-converters to AD). Introducing machine learning to boost performance in a data-driven way may enhance the identification of brain regions that are known to represent

Alzheimer’s Disease. This technique is commonly found in 2D MRI slices to identify the most meaningful slices that detect areas of the brain known to represent Alzheimer’s Disease.

An intuitive approach used in [16] combines a 2D CNN with Regions of Interest patches (ROI) known to represent AD. These local patches are assembled into 2.5 dimensional patches to train the CNN by identifying informative features in multiple regions of the brain such as the hippocampus and temporal lobe. They hypothesize that by extracting local patches, more high-level features can be identified by extracting areas of the brain known to represent Alzheimer’s Disease. Finally, the trained CNN-based features are combined into a feed-forward machine learning classifier to identify MCI converters to Alzheimer’s Disease. This is done by calculating the outputs as a formula to avoid the random generation of an input weight matrix, thus classifying as either converter/non-converters using these features [16]. They report an accuracy of 68.49% in the classification of converters/non-converters when trained with converters/non-converters patches. Low performance accuracies in MCI converters are common due to less degeneration in brain tissue that can be harder to identify compared to a patient who has converted to Alzheimer’s Disease. However, they report a higher accuracy of 73.04% when trained with AD/HC patches. This tends to be a common trend amongst researchers in an effort to diagnose the early stages of mild cognitive impairment where identification of discriminable changes is much more subtle.

2.3 Methods

In this chapter, we discuss the dataset used in both training and testing of our model and the processing pipeline that was required. MRI scans require multiple preprocessing steps that are crucial to training because of the image disparity than can be found between patients.

2.3.1 Dataset

To evaluate the feasibility of our model, magnetic resonance imaging (MRI) scans were downloaded from the Alzheimer's Disease Neuroimaging Initiative. Founded in 2003 by lead investor Dr. Michael W. Weiner, ADNI has become a leader in the field of neuroscience research, with a focus on evaluating the progression of mild cognitive impairment (MCI) and Alzheimer's disease (AD) through MRI and positron emission tomography (PET) scans, biomarkers, and neuropsychological evaluations. The ADNI 1.5T collection was downloaded because of its preprocessing pipeline including 3D gradwarp and B1 non-uniformity correction to prepare the scans for further processing. Consisting of three-axis slices (axial, coronal, sagittal), a 3D image is created by combining all planes. A total of 750 three-dimensional scans were used and divided evenly between Alzheimer's Disease patients and a cognitively normal control group, respectively (n = 750 scans: AD = 375, HC = 375).

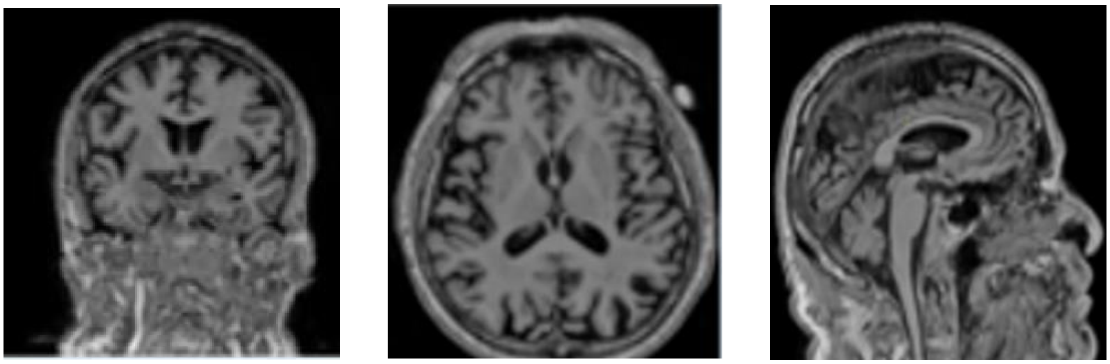

Figure 2.2: cross-section of brain, consisting of the coronal (left), axial (middle), and sagittal (right) slices.

The average age of participants was 75y, comprising of T1 scans in the neuroimaging informatics technology initiative (NIFTI, .nii) format. T1 scans were used to brighten tissues with a high-fat content (white matter), allowing them to be more visible in the scan, while watery gray matter structures tend to appear darker [17]. The detection of degenerative brain tissue throughout multiple regions of interest (ROI) such as the hippocampus and cerebral cortex offers a prospective indication of AD. This is done by utilizing volumetric data to obtain full brain mapping within an MRI. This collection of data from ADNI is commonly used in AD research and can be seen in [5] to facilitate training.

2.3.2 MRI Preprocessing

Preprocessing of the MRI data was carried out using the CAT12 Toolkit pipeline with default settings. The pipeline includes affine regularization, spatial normalization, MNI space registration, image smoothing, white matter (WM) segmentation, and skull extraction on all images. Cat12 maps the brain to yield a 3D voxel space representing a specific volume within the image of (X, Y, Z) dimensions according to height, width, and depth. This creates volumetric three-dimensional data in the same process as [17]. After processing, each image had a new dimension of 121 x 145 x 121 voxels, with a spatial resolution of 1.5 x 1.5 x 1.5 mm^3 for each voxel. The pipeline used in this research follows the same procedure as [17] for 3D brain segmentation through the Cat12 toolkit.

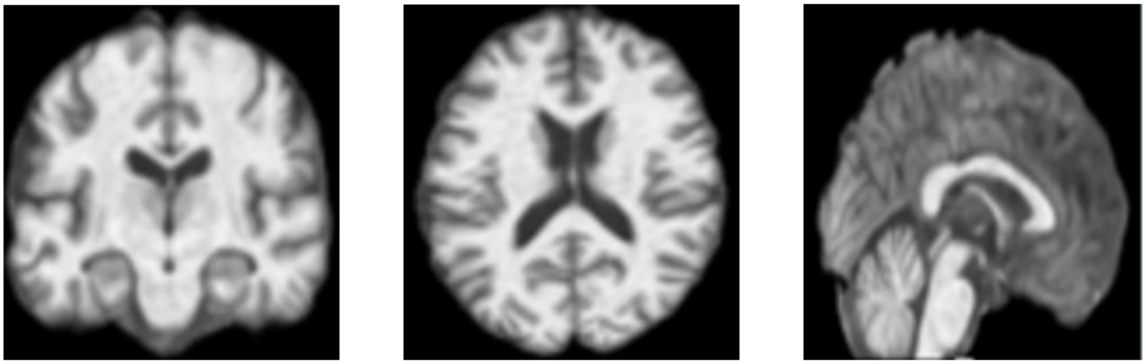

Figure 2.3: Post-processed MRI images after CAT12 Toolkit.

The key importance to understanding neurodegeneration within 3D MRI scans falls within the domain of voxel-based morphometry (VBN), calculating the size of an individual voxel to compare with others in the local area. The comparison of voxels-values may offer distinguishable insights to volumetric changes between subjects of different groups. Understanding these changes allows for the comparison and surveillance of neurodegeneration over a period of time, particularly in MCI and AD patients. VBM performs voxel-wise comparisons of gray matter concentration between two groups of subjects, in our case being NC vs AD. This will allow for the observation of volumetric changes within areas of the brain known to represent the advancement of AD. The following steps display the processing pipeline used by Cat12 to perform VBM analysis:

1. Image Segmentation: Segmentation was achieved through white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF) tissue probability maps to identify and highlight the different matter of the brain. In our case, WM illumination is critical because it houses nerve fibers, known as axons, which carry electrical signals to other neurons and can show degeneration linked to AD.
2. Spatial Normalization: Due to variations in skull size, head position, and orientation, the alignment of patients is necessary to provide spatial normalization to all images being processed. By registering each image to the standard Montreal Neurological Initiative

(MNI) space, affine regularization is used to perform linear mapping and achieve global geometric transformation for all images. Lastly, to improve resolution and quality, a non-linear transformation is done to minimize the regional differences in local deformations.

3. Smoothing: A Gaussian filter is then applied to the segmented images to "clean up" the warping effects caused by spatial normalization which improves the signal ratio, thus increasing statistical power.
4. Skull Stripping: The last processing step includes skull stripping to create a mask of the brain by removing all skull and dura matter. Cat12 uses an adaptive probability region-growing (APRG) method to combine both GM and WM probability maps to yield a bias corrected mask of the scan

Figure 2.4: Data augmentation is performed through random rotations within the 3D image.

Deep learning model performance has shown promising results in image classification tasks due to their ability to recognize features. Because of this stipulation, model performance is predicated on the characteristics of the dataset including the number of samples and the processing techniques that may have already been performed. Data augmentation is an important step to defend against overfitting of the model when there is low disparity between images. Particularly in medical imaging, it can be difficult to obtain large datasets due to limited online resources, patient privacy rights, and the amount of labor needed to label and organize data. This allows smaller datasets to function as larger ones due to the manipulation and "augmentation" of each image to promote disparity and emulate a larger volume. In this study, we have obtained a semi-large dataset of 750 images that have been flipped, rotated, and gaussian blurred to perform image augmentation and reduce any imbalances in the dataset.

2.4 Deep Learning Approach

The CNN has displayed tremendous success as an image recognition tool that specializes in feature detection and can outperform its predecessors such as feedback neural networks by detecting learnable features within an image. In recent studies, 3D image classification has shown promising performance of AD classification through generic feature segmentation while using a CNN to reduce complexity [26]. The ability to detect features and inform a proper diagnosis has become integral in medical imaging to detect abnormalities within radiology scans such as MRI and PET [11]. Efficient diagnosis can offer alternative treatment plans such as medications and therapy that are typically accompanied by early detection of cancerous cells and degenerative tissues. The CNN will pass the MRI scans through multiple convolutional layers including filtering layers, pooling layers, and connected layers to provide a probabilistic output between 0 and 1 using the sigmoid function.

2.4.1 Convolutional Neural Networks

Convolutional neural networks (CNN) boast the unique ability to accept images as direct input to perform calculations to learn local features, compared to other methods of machine learning that rely on raw text values to perform techniques such as classification and regression. Adding dimensionality to a model further increases the learning potential by introducing more features that can be captured within an image and then being vectorized before training. In our case, we aim to test the efficacy of 3D networks that take advantage of the third spatial dimension to optimize variance and achieve predictability within our model.

2.4.2 2D vs. 3D Networks

Studies within the domain typically rely on the use of 2D networks to detect degenerative changes in the brain by slicing the image into cross-sectional planes. While there are several studies that use 3D networks for classification, 2D networks remain traditional due to their more simplistic architecture with less complexity in learning parameters and image size. However, while this may be beneficial for simplistic images with less features, a 2D approach may not be as suitable for medical imaging due to its inability to capture subtle changes while not utilizing the third spatial dimension. Because of this, most 2D networks are accompanied by multiple learning techniques including ensemble learning [17] and the use of ROI patches [16].

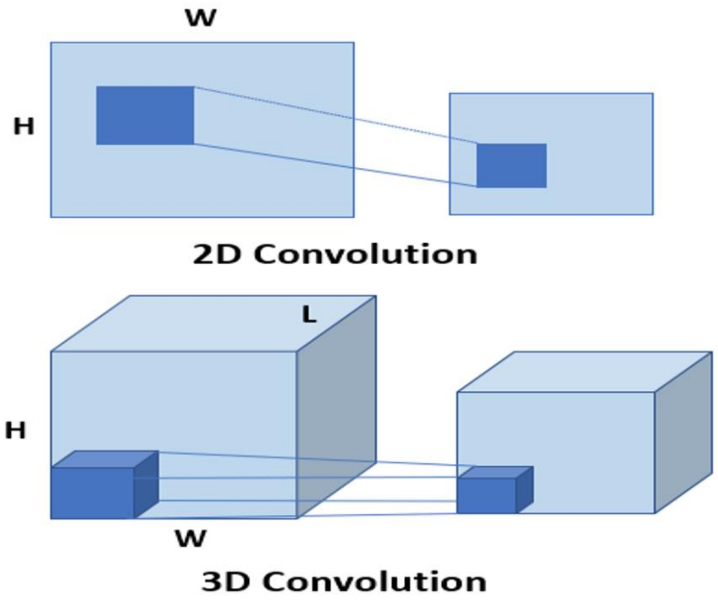

Figure 2.5: 2D and 3D convolutions map areas of an image to learn important features.

When analyzing medical images, it can be at times quite difficult to “eye-ball” physical changes within the image, depending on the nature of scan, and make an accurate assumption of how much change a tissue or mass has overgone. In the case of monitoring MCI, it’s imperative to understanding the timeline and progression of AD to inform a proper diagnosis. This idea displays the need for volumetric data when attempting to detect degenerative tissues and changes within brain volumes. Taking advantage of the third spatial dimension creates an opportunity to capture these changes by observing not only the height and width, but depth an area. By using 3D convolutional layers, important local patterns can be learned in a depth-wise approach by introducing complexity in the model directly correlated by the specified number of filters in a convolutional layer. This introduces the ability to learn even deeper features within an image to influence successful classification between multiple classes. While introducing complexity adds high variance, it also decreases the bias of your model because of the inability to generalize between images due to extreme model complexity, leading to overfitting. Because 3D architectures tend to be more complex, it’s imperative to understand the median between model variance and bias, verifying that your model has predictive capability while reducing generalization.

2.4.3 Feature Extraction

The CNN utilizes the connectiveness of neural networks to perform image analysis by applying convolutional layers with a list of filters. The number of filters denotes the filter mask size of the convolving window that moves around the image to detect different shapes and edges. Increasing the filter size in subsequent layers may allow the CNN to capture more complex features due to the larger area of pixel coverage. It is important to note that the number of kernel parameters change depending on the dimensionality of the network, meaning that a 3D network will have a filter size of (3 x 3 x 3 pixels) compared to a 2D network with dimensions (3 x 3 pixels). Thus, ignoring the third dimension will exclude the depth of the scan where features may be more detectable utilizing a volumetric approach. The CNN performs well in feature extraction due to its “deep” architecture by stacking convolutional layers and down-sampling to increase the area view for each neuron.

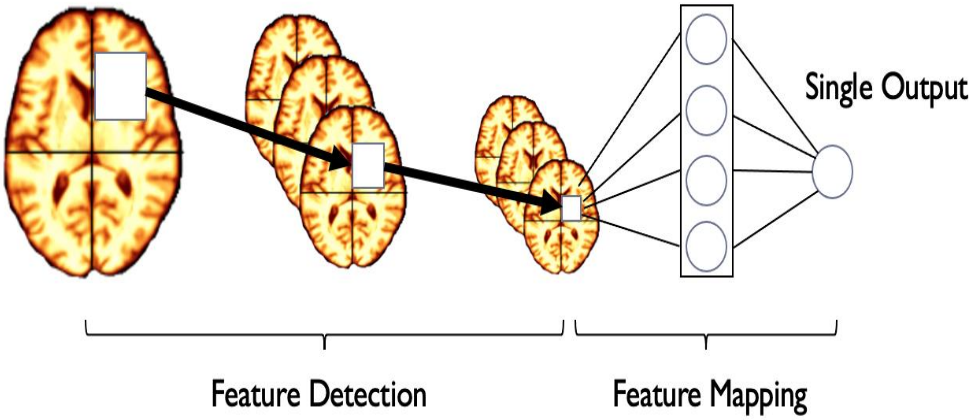

Figure 2.6: 2D CNN performing feature detection and mapping.

CNN’s can be divided into two main functions: feature detection and feature extracting are used to obtain learnable parameters by calculating adjacent spatial dimensions (X, Y, X) within an image [26]. Learnable filters (i.e., weights) are applied in feature detection where each neuron is connected through receptive neurons of the previous layer to extract low-level features within an image [26]. Feature extraction is then performed by applying convolutional filters to the input to create feature maps of local areas. The down-sampling of spatial dimensions allow for max-pooling, batch normalization, and fully connected layers to increase the receptive field of each neuron to learn low level features. Lastly, the combining of all neurons into a single dense output of one allows the CNN to effectively reduce the number of parameters in the network as well as the reduction of feature maps into a binary prediction between two classes.

2.4.4 Model Architecture

Our proposed 3D-CNN architecture consists of five sequential layers of conv3D (with Relu activation function), batch normalization, and max pooling layers. Sequential convolutional layers with local filters of 3 x 3 pixels were used to extract learnable parameters within MRI scans to capture spatial and positional relationships throughout the brain. Each iteration of conv3D applies 3-dimensional filtering (X, Y, Z) among the axial, coronal, and sagittal planes. Following is batch normalization, which stabilizes image weights and standardizes the inputs, and max pooling which down-samples the feature maps to iteratively reduce the image size. The rectified linear unit (ReLU) activation function was used in each convolutional layer to apply a linear identity to our input by returning 0 representing a negative input and 1 as a positive output. Below is the network architecture:

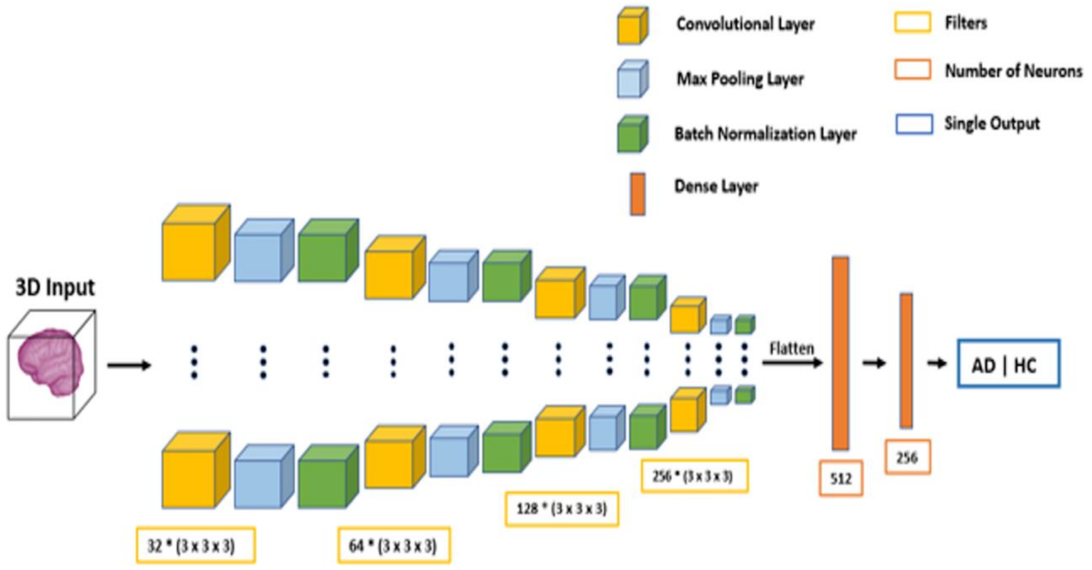

Figure 2.7: The complete network topology

Following the four sequential layers, global average pooling is used to calculate the average output for the feature maps of previous layers, concluding with a fully connected dense layer of 512 to flatten all neurons in the network. Lastly, a dense layer of 1 is used to combine all neurons into a final connected output of 1. For the loss function, binary cross-entropy is used to yield a binary classification between two classes of HC and AD. Lastly, the sigmoid function is applied to find the probabilistic output of the model between 0 and 1 through binary classification. Dropout was used after the first and third convolutional layers as well as before and after the fully connected layers to reduce the risk of overfitting by randomly setting neurons in the previous layer to 0 after each iteration. Experiments were performed on a single node RTX 2060 Super GPU. Model performance was tested for 200 epochs with an average epoch computation time of 132 seconds, for an average total runtime of 8 hours. 7

2.5 Results

A classification experiment between two classes, Alzheimer's disease (AD) patients and a healthy cognitive control group (HC) was performed to classify MRI scans as either HC or AD through binary classification. Consisting of 750 scans, an 80-20 train-test split was performed to utilize a training set for the model to learn features and a test dataset to perform validation against the train set. To avoid data leakage and promote image disparity, there was no overlap in scans between train and test classes. Binary classification was then performed between HC and AD, yielding an accuracy of 79%. Results of classification between HC and AD have shown to be at par or comparable to leading papers in this research such as [17], with an accuracy of .84 ± .03 within a 2D CNN accompanied by ensemble learning (EL).

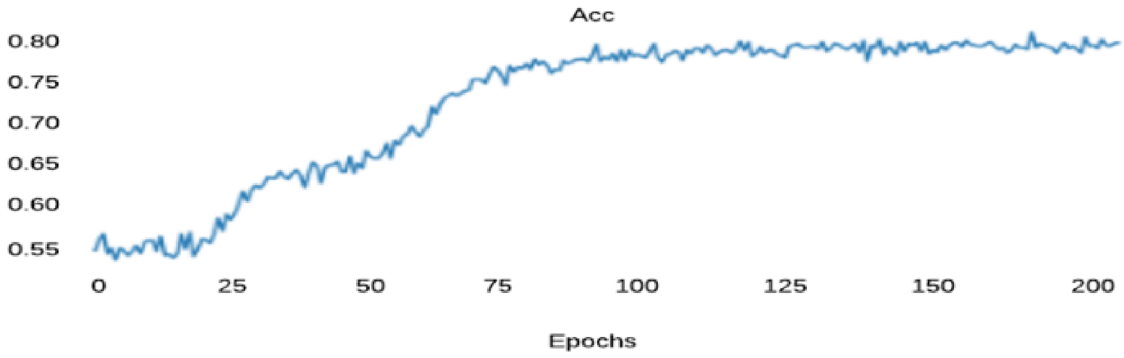

Figure 2.8: Model accuracy during training over 200 epochs.

The goal of this research was to show the efficacy of 3D convolutional networks in classification tasks while taking advantage of volumetric data to detect abnormalities. The results achieved offer insight into 3D models compared to other techniques such as support vector machine (SVM), random forest (RF), and principal component analysis (PCA) for classification tasks. The current results of our model show an accuracy that slightly outperforms the following study [5] when using SVM and PCA for classification tasks in MRI scans for AD detection (76% acc). To save time and computing costs, 2D models are also popular in classification tasks that take 2D images of either the axial, coronal, or sagittal planes to detect abnormalities within the angle of each plane. While being a simpler approach, 2D networks lack in extracting spatial information in non-volumetric data that is needed to efficiently map degeneration [25].

2.6 Streamlit Application

The effort to demonstrate this research has been achieved by developing a web-application to return probabilistic predictions on test images between NC and AD. The development of a web-application allows for the demonstration of our model by predicting on single test images and displays the three cross-sectional slices of the brain. The idea for this application was centered around simplicity in its deployment and reliability after launch. To solve these problems, we utilized Streamlit, an open-source Python library used for launching machine learning applications both quickly and efficiently. Built off Python, Streamlit allows for the integration of many different libraries and API's to quickly develop ML applications for either demonstration or business purposes. The Streamlit library allows for the simple creation of widgets and templates to display data, accept input, and utilize your own model for testing purposes. This is extremely beneficial for professionals in the domain who are not concerned with web development and are only concerned with deploying their work. Thus, Streamlit was used to develop a simple webpage structure to demonstration our model by classifying MRI images between NC and AD patients.

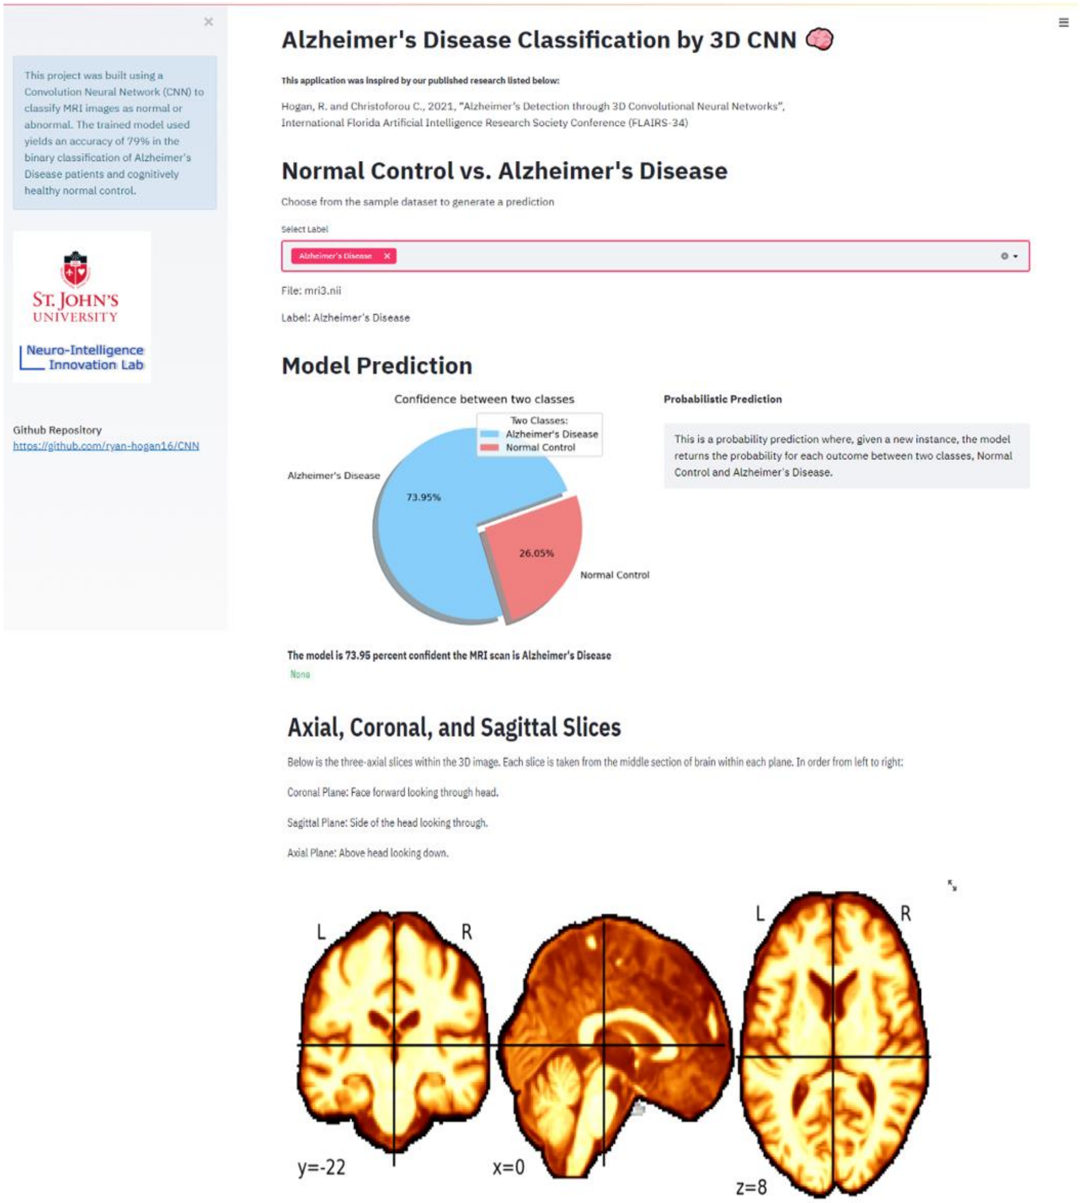

This project was built using a Convolution Neural Network (CNN) to classify MRI images as normal or abnormal. The trained model used yields an accuracy of 79% in the binary classification of Alzheimer's Disease patients and cognitively healthy normal control.

Github Repository
https://github.com/ryan-hogan16/CNN

Alzheimer's Disease Classification by 3D CNN

This application was inspired by our published research listed below:

Hogan, R. and Christoforou C., 2021, "Alzheimer's Detection through 3D Convolutional Neural Networks", International Florida Artificial Intelligence Research Society Conference (FLAIRS-34)

Normal Control vs. Alzheimer's Disease

Choose from the sample dataset to generate a prediction

Select Label: Alzheimer's Disease

File: mri3.nii

Label: Alzheimer's Disease

Model Prediction

Probabilistic Prediction

This is a probability prediction where, given a new instance, the model returns the probability for each outcome between two classes, Normal Control and Alzheimer's Disease.

The model is 73.95 percent confident the MRI scan is Alzheimer's Disease

None

Axial, Coronal, and Sagittal Slices

Below is the three-axial slices within the 3D image. Each slice is taken from the middle section of brain within each plane. In order from left to right:

Coronal Plane: Face forward looking through head.

Sagittal Plane: Side of the head looking through.

Axial Plane: Above head looking down.

Figure 2.9: Example output of the web application

2.7 Discussion

To inform a reputable classification with 3D images while using a CNN, this study is broken up into four key components: (1) Data collection was done by downloading T1 weighted MRI scans from ADNI to facilitate testing between both HC and AD patients. The data was downloaded as 3D images in the NIFIT file format, which is commonly used in neuroimaging research and follows the same collection used in [5]. Upon downloading, the images had already been processed through the ADNI pipeline to perform image correction and scaling. The images used consisted of cognitively healthy control patients as well as Alzheimer's disease patients who have converted from MCI.

(2) MRI preprocessing was performed to normalize all images including white matter segmentation to highlight WM which exhibits degeneration in the brain. Preprocessing included registration to the MNI space, affine registration, spatial normalization, image smoothing, and image resizing of 121 x 145 x 121. These steps allow for the normalization of the brain volume to be within a range of 0 and 1. Image preprocessing is necessary for successful classification tasks and followed the same procedures seen in [17].

(3) The CNN was used to perform feature extraction and learn spatial information within adjacent pixels. This allows for the reduction of learnable parameters and has been a staple of image classification within medical imaging [24]. A four-layer convolutional stack was used including sequential convolutional layers, pooling layers, and batch normalization layers to reduce features and derive a binary classification of 1 between two classes HC and AD.

(4) An average binary classification accuracy of 0.79 was achieved using a 3D architecture to detect abnormalities within volumetric data. Results of this method can be comparable to other leading studies such as [5] who achieved an accuracy of .76 while using a 3D architecture including ROI patches.

CHAPTER 3

ECG CLASSIFICATION USING LONG SHORT-TERM MEMORY NETWORKS

3.1 Introduction to Heart Arrhythmias

Cardiovascular diseases and other underlying heart conditions continue to damage our society as being the number one cause of death in America, claiming a life every 36 seconds [12]. Heart conditions such as arrhythmias, hypertension, and strokes can be life-threatening to individuals who maintain a poor diet and carry a sedentary lifestyle. This has proven to be an issue within the United States where only one in four adults maintain an adequately healthy lifestyle requiring enough aerobic and muscular-strengthening activities to reduce the risk of cardiovascular diseases and others alike [12]. However, not all cardiovascular issues stem from lifestyle choices such as poor diet, smoking, and lack of physical activity. Heart arrhythmias present life-threatening conditions due to irregularities in heart function that may be predisposed or can develop throughout a lifetime. These irregularities in heart function can affect an individual in numerous ways including dizziness, irregular beat patterns, and rapid heart rate that depending on the severity of the arrhythmia, and the predisposal of the individual due to genetics and lifestyle choices, may result in life-threatening consequences.

The intuition of this research stems from personal experience, while I suffered from an underlying heart condition called Wolf Parkinson White Syndrome (WPW). WPW syndrome is a congenital heart condition where an extra electrical pathway between the upper and lower chambers of the heart causes irregular and elevated heart rates, typically accompanied by exercise and is usually detected during adolescence. My condition went undetected for over a year despite the use of a heart monitor at all times, even during athletic activity. When finally detected by ECG, my heart rate exceeded 220 BPM, a rate at which could be fatal to individuals with previous health concerns or are of older age. Future research will include the detection of WPW in ECG recordings by monitoring the preexcitation period of the heartbeat, including the potential signs that denote abnormal heart function that an individual may have and not be aware of.

Arrhythmias may not always be prevalent to an induvial who shows no symptoms and may be considered a healthy individual. The detection of arrhythmias can go unnoticed because of the timing of the arrhythmias as well as the conditions, such as strenuous activity, that may be necessary to spike such occurrence. To monitor arrhythmias and heart functionality, an electrocardiogram (ECG) is the industry standard in identifying the rhythmic changes of the heart while detecting abnormalities within beat patterns through each cardiac cycle. The use of ECG monitors the cellular depolarization and repolarization of the heartbeat within a period of time, typically a ten-second strip [9]. The detection of abnormalities is prudent to proper treatment options depending on the severity and type of arrhythmia present.

The human body is considered to be a large conductor of electrical signals where conductive ionic fluid in the body allows for recording the signals at the skin-level [9]. The importance of detecting these signals allow cardiologists to examine and diagnose potential abnormalities within beat patterns. Because ECG is time-series data, the specific electrical signals within the beat pattern can be identified to determine if there are any potential complications within the heart cycle.

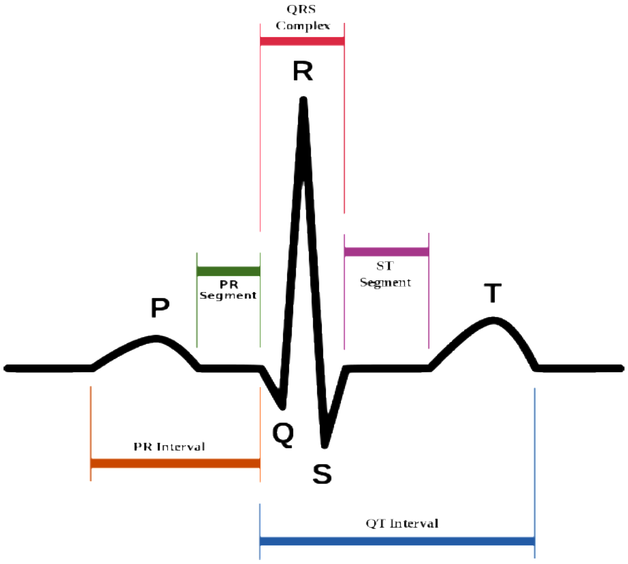

Figure 3.1: A single ECG wave (Wikimedia Foundation, 2020).

The goal of this research was to offer an alternate form of arrhythmia detection in ECG recording through the use of deep learning and LSTM neural networks. In our network, we will perform binary classification of normal heartbeats, and a superclass of eight arrhythmias including supraventricular ectopic beat, ventricular beat, fusion beat, myocardial infarction, myocarditis, valvular heart disease, myocardial hypertrophy, and bundle branch block.

LSTM, as well as other deep learning methods have shown promising results in both binary and multi-class classification of normal and abnormal heartbeats within raw ECG data as seen in [13], achieving an accuracy of 95% in the classification of ECG signals using a Convolutional Neural Networks (CNN) + LSTM. It is necessary to note that while LSTM performs very well at classifying time-series data, other techniques have also shown promising results [23] by utilizing 2D CNN’s that take ECG images as input to identify discriminable changes between normal and abnormal beats with an accuracy of 97%.

3.2 Related Work

This paper aims to identify the classification of abnormal heartbeats through the use of long short-term memory (LSTM) networks that excel at recording vital information in time-series data by forgetting non-informative features during training. We propose a two-layer network of 50 and 100 units, respectively, to provide a reputable performance compared to preceding research in the domain.

In recent years, a deep learning approach has been applied to identify abnormal arrhythmias in ECG recordings that are directly linked to both cardiac and non-cardiac disorders. Research in this field attempts to classify arrhythmias as either a binary classification, being a either normal beat, or abnormal beat considered. The abnormal class is made up of a super-class of arrhythmias for multi-class classification problems, where specific beats can be categorized within the dataset according to diagnostic class. The approach used by [25] performed binary classification through multiple techniques such as a recurrent neural network (RNN), RNN + gated recurrent neural network (GRU), and RNN + LSTM to identify heartbeats as either normal or abnormal. They report classification accuracies of 85.4%, 82.5%, and 88.1%, respectively, for RNN, RNN + GRU, and RNN + LSTM.

The results show that LSTM tends to report better results than a vanilla RNN network due to the effectiveness of using multiple hidden layers within the network. The dataset used was from the MIT-BIH Arrhythmia dataset, which is commonly compared to the dataset in our work from the PTB Diagnostic ECG Database and is often interchanged in this field of research. We hypothesize that [25] could improve performance by fine-tuning the number of activation units in the network's hidden layers to learn more informative features in the ECG recordings. The study proposed used activation units of 64, 100, and 256, respectively, whereas we used 50 and 100 units in our hidden layers. The dropout rate can also be tuned to improve performance, where a dropout of 0.2 is used in every layer to drop neurons in the network to reduce overfitting and optimize the weights in the network.

An alternate deep learning method can also be used to detect abnormal heartbeats in ECG recording with a convolutional neural network (CNN). The CNN continues to be a prominent form of feature detection in classification tasks by accepting images as input to identify features in contiguous local areas. The CNN approach was used by [27] to perform binary classification of ECG recordings that have been preprocessed using the logarithmic transformation of the time-series data. After processing, the ECG recordings were fed into the CNN, consisting of 6 convolutional blocks with filters of 5 x 5. Temporal averaging is used to process the variable-length input ECG signals before they are fed into the linear SoftMax classifier, yielding a label of a probabilistic output between two classes. An accuracy of 90.5% was reported when only using the CNN to predict between normal and abnormal arrhythmias, which is at par with current papers in the domain. The CNN may not be suitable for time-series data due to potential information leakage because of increased parameters in the network and the need for aggressive data augmentation such as logarithmic transformation.

3.3 Methods

To evaluate the feasibility of our model, the Physikalisch-Techhnische Bundesanstalt (PTB) Diagnostic ECG Database [1] was used to classify between normal and abnormal arrhythmias. The collection of ECG recordings was pioneered by Professor

Michael Oeff, M.D., at the Department of Cardiology of University Clinic Benjamin Franklin in Berlin, Germany. The use of this dataset will help facilitate future discoveries in the cardiology field with its extensive amount of data as well as the necessary preprocessing steps to facilitate testing.

3.3.1 Dataset

The PTB database utilizes 16 input channels (14 for ECGs, 1 for respiration, and 1 for line voltage), with an input voltage of ±16 mV, with a compensated offset voltage up to ± 300 mV. Each ECG signal is digitized at 1000 samples per second with a 16-bit resolution and a sampling frequency of 125Hz. The database contains 549 records from 290 subjects with an age range of 17 to 87, while each subject contains 15 ECG signals, respectively. Below are the summary statistics of the participants of the database, as well as the diagnostic class for subjects within the database (the diagnostic class of 22 subjects was not available).

	N	**Mean Age**
M	209	57.2
F	81	55.5

Table 3.1: Summary statistics for participants in the PTB Database.

Of the diagnostic class, myocardial infarction leads in the number of subjects by a large margin. Myocardial infarction is a result of a blot clot in the epicardial artery that supplies blood to the heart, otherwise known as a heart attack [22]. Heart attacks remain one of the leading cardiovascular diseases in adults, so it does not seem uncommon that myocardial infarction makes up 71% of the diagnostic class.

Diagnostic Class	Number of Subjects
Myocardial Infarction	148
Cardiomyopathy	18
Bundle Branch Block	15
Dysrhythmia	14
Myocardial Hypertrophy	7
Valvular Heart Disease	6
Myocarditis	4
Miscellaneous	4
Normal Control	52

Table 3.2: Subject breakdown according to diagnostic class

3.3.2 ECG Preprocessing

Of the diagnostic class, myocardial infarction leads in the number of subjects by a large margin. Myocardial infarction is a result of a blot clot in the epicardial artery that supplies blood to the heart, otherwise known as a heart attack [22]. Heart attacks remain one of the leading cardiovascular diseases in adults, so it does not seem uncommon that myocardial infarction makes up 71% of the diagnostic class.

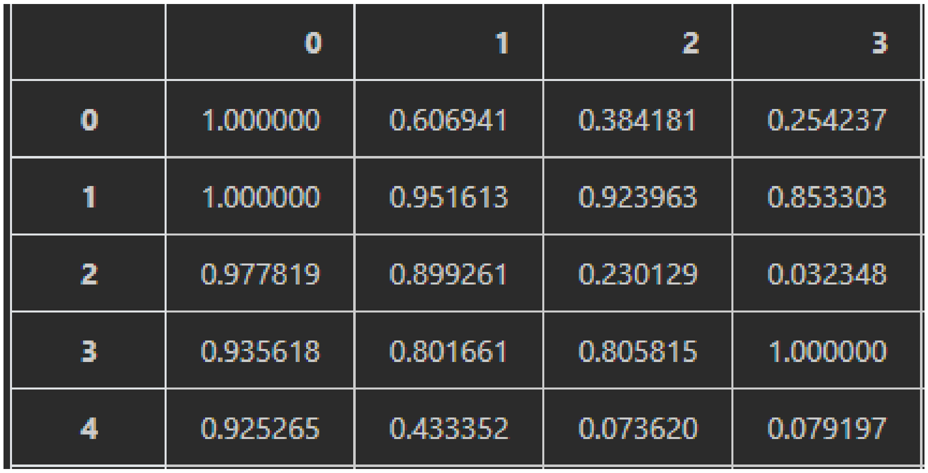

	0	1	2	3
0	1.000000	0.606941	0.384181	0.254237
1	1.000000	0.951613	0.923963	0.853303
2	0.977819	0.899261	0.230129	0.032348
3	0.935618	0.801661	0.805815	1.000000
4	0.925265	0.433352	0.073620	0.079197

Figure 3.2: The time-series data of an ECG recording

All necessary preprocessing steps have already been performed on acquisition of the dataset. When constructing the train-test split, the label of each row was assigned as the target label 'y' to classify between normal and abnormal recordings. Below is example of a single plotted ECG recording.

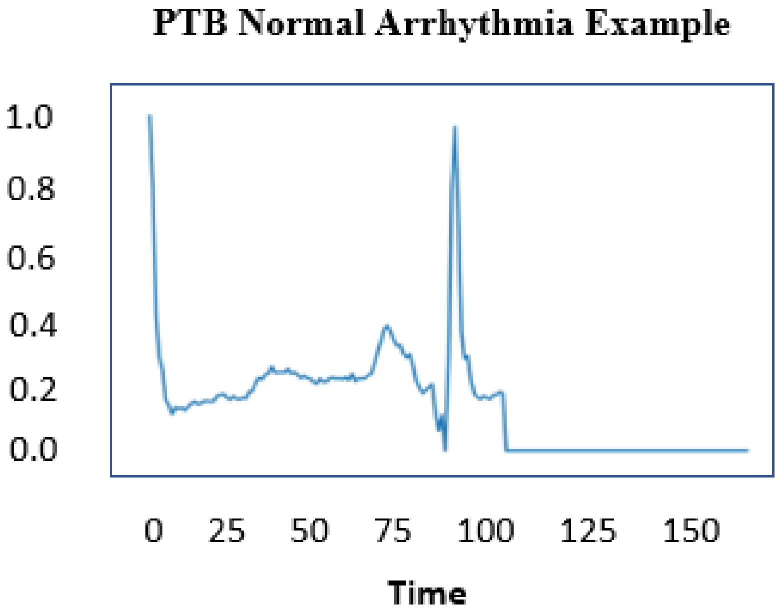

Figure 3.3: A single ECG recording that is plotted from a row of the PTB dataset.

3.4 Long Short-Term Memory Networks

In order to capture the subtle changes in time series data, LSTM utilizes *memory blocks* in the hidden layer to remember and forget various sequences to form a prediction [21]. These memory blocks control the mechanism of the network by storing the temporal state of the memory cells while being connected to *input* and *output* gates. Memory blocks perform well at remembering information throughout the loop of the network, allowing information to persist throughout the hidden layers and updated cell states. LSTM consists of four neural network layers that are used to store information from previous states and use two activation functions (tanh & sigmoid) to regularize both inputs and outputs. The LSTM repeating architecture is shown below.

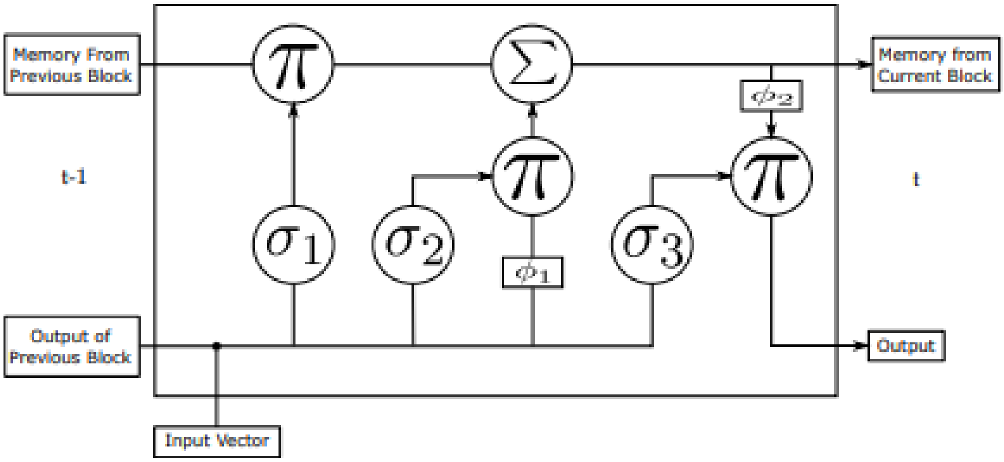

Figure 3.4: The generic LSTM architecture consists of four hidden layers [15]

The first process in LSTM is through the initial forget gate layer using the sigmoid function to output a number between 0 and 1. The closer the output value is to 0 is considered not valuable information, while closer to 1 being extremely valuable information for the network to remember. This initial mechanism allows the network to understand parameters when given an input. The cell state of the previous input may need to be updated relative to the new information being passed, such as a different type of arrhythmia that was determined in the previous input [15].

$$cf_1 = \sigma_1(W_{cf} \cdot [O_{t-1}, x_t] + b_{cf})$$

The next process is done in two steps, where the sigmoid function acts as the input gate to decide which values are necessary to update and the tanh function creates a vector of these new candidate values [15].

$$I_t = \sigma 2(W_I \cdot [O_{t-1}, xt] + b_I)$$

The next layer updates the old cell state into a new state *St* by using the sigmoid layer ($\sigma 2$) which decides the values to update for the next layer, and tanh layer $\phi 1$ to derive the new candidate values [15].

$$S_t = tanh(W_S \cdot [O_{t-1}, xt] + b_S)$$

Lastly, the cell state is updated using the function below, creating a final updated cell state as well as an output from the network. This cell state will be used for the next iteration of the network and may or may not be updated depending on the input [15].

$$S_t = cf \cdot S_{t1} + I_t \cdot S_{t-1}$$

3.4.1 LSTM Architecture

The proposed network architecture utilizes a simple LSTM scheme to identify arrhythmias within ECG recordings. Identifying the temporal sequences within the ECG scan allows LSTM to develop long-range dependencies within the wavelengths, forming a prediction between two classes normal and abnormal.

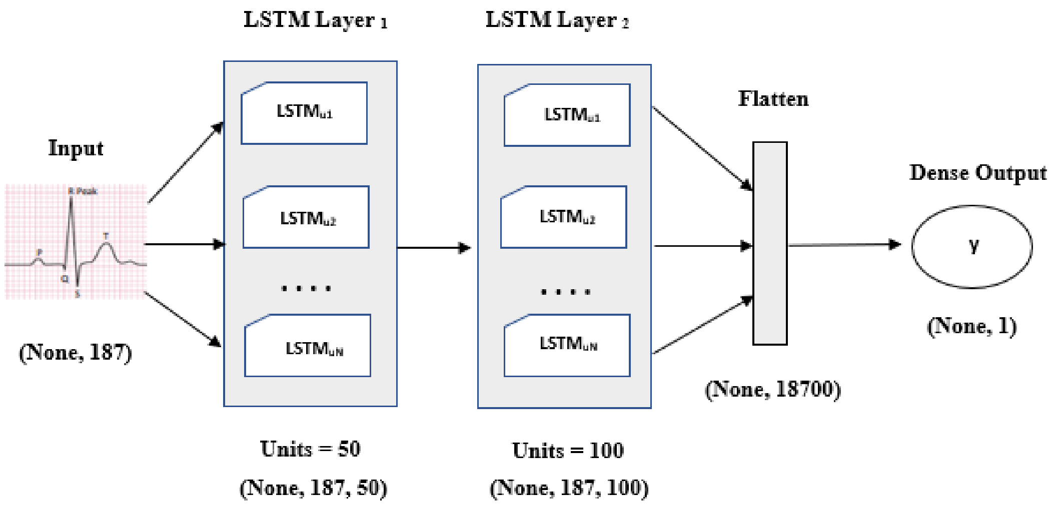

Figure 3.5: Custom LSMT architecture used.

The LSTM architecture begins by taking a one-dimensional ECG recording as input, passing it to the two LSTM layers. The model contains two LSTM layers of 50 and 100 units, which determines the dimensionality of the hidden state and will return a sequence vector of dimension 100. Regularization was also used in a single dropout layer of (0.2) after the first LSTM layer to stabilize weights within the network. The network is then flattened into a single vector, expanding to the size of (None, 18700). Lastly, the sigmoid function is used to derive a single dense output through binary classification. The output of the model will be a probabilistic prediction of either a normal ECG recording or an abnormal recording. A batch size of 128 was used over 30 epochs during training. A training set of 10,185 samples were used to train our model, and 4,365 to validate it. Training was performed on a single-node RTX 2070 Super GPU.

3.5 Results

The performance of our model was evaluated through metrics such as accuracy, precision, and recall understanding the predictions our model is making. The goal of model performance is to predict True Positives and eliminate the prediction of True Negatives, or even False Positives. Accuracy refers to the total number of correct classifications between normal and abnormal beat patterns ((TP + TN) / (TP + TN + FP + FN)) . Precision is also used to measure the number of True Positives divided by the number of True Positives plus the number of False Positives (TP / (TP + TN)). Lastly, recall is considered to be the True Positive rate of the model, making the most relevant predictions (TP / (TP + FN)).

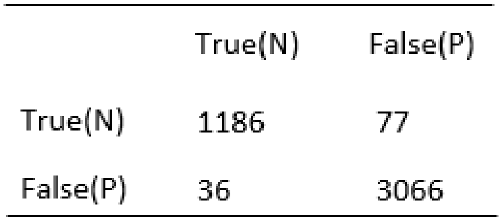

	True(N)	False(P)
True(N)	1186	77
False(P)	36	3066

Table 3.3: Confusion matrix

A classification experiment between two classes was formed to identify arrhythmias within ECG recordings. The proposed LSTM model boasts an accuracy performance of up to 97% in the binary classification of both normal and abnormal heartbeats. The model performed extremely well in classifying ECG recordings due to LSTM's ability to recognize temporal features in time-series data. Below is the complete table of model performance in ECG classification:

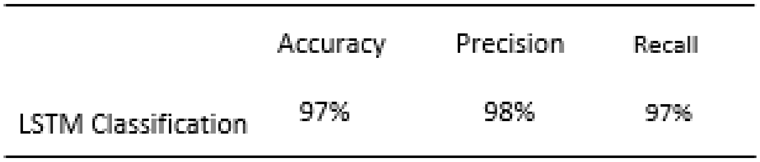

	Accuracy	Precision	Recall
LSTM Classification	97%	98%	97%

Table 3.4: Model performance for accuracy, precision, and recall

	Approach	**Accuracy**
Our Approach	LSTM	97%
Singh et al., 2018	LSTM	88%
Klosowski et al., 2018	CNN + LSTM	86%
Zihlmann et al., 2017	CNN + LSTM	82%

Table 3.5: Performance accuracies of other models in ECG detection

The detection of arrhythmias in ECG recordings using deep learning has been a relatively new practice in the research community. While there are multiple ways of detecting such arrhythmias including Convolutional Neural Networks, Stacked Auto-Encoders, and LSTM, all offer different forms of detection using neural networks. The data being used may influence the specific network topology and approach in arrhythmia detection.

CHAPTER 4

CONCLUSION

4.1 Future Research

In this study, we introduce two computational approaches to address real-world medical diagnosis problems that continue to burden our society. The use of deep learning models has offered an alternate form of detection and understanding of health conditions in a data-driven way. This ushers in a new area of medical research that involves a potentially more accurate method of diagnosis compared to leading practices. Technology has reached a point where both the availability and amount of data can finally be harvested and applied artificial learning techniques, particularly in the healthcare domain. Research in this field continues to grow and will facilitate discoveries in healthcare as well as promote a better understanding of health conditions alike.

An important task that is not currently presented in this study is the early detection of MCI to AD within MRI images. Mild cognitive impairment is an important turning point in the progression of AD when configuring treatment remedies and health care options. The need for early classification requires more attention to detail due to the subtle changes in the brain during mild impairment compared to the changes seen in Alzheimer's Disease patients. Classification accuracies between MCI and AD have not shown to outperform NC and AD when detecting neurodegeneration due to the less extreme changes within the brain volume. The following study reached an accuracy of 62% when testing MCI vs AD, and 84% for NC vs AD which is a considerable difference [17].

Efforts are being made to create a model to classify between MCI and AD with an accuracy rate that approaches or exceeds current accuracies in literature. We have attempted to train our model to classify between MCI and AD with an accuracy of 58%. With an accuracy of barely over half, results were not included for this study with efforts of raising accuracies scores at par with the industry standards. Several considerations can be made including the fine-tuning of model parameters, increasing the data size, and transferring learning to potentially increase scores to reputable heights.

The purpose to detect and classify heart arrhythmias is a personal endeavor that has been directly experienced, motivating the purpose of this research. Heart arrhythmias may go unnoticed to seemingly healthy individuals that present no symptoms and carry a healthy lifestyle. Future research in the detection of heart arrhythmias will include the screening of WPW in the ECG time-series data by identifying time-interval changes in the QRS complex of the heartbeat, a segment known for representing preexcitation in the heart cycle. My condition went undetected for over a year, despite the constant use of a heart monitor designed to trigger when an abnormal arrhythmia had presented itself. However, WPW is not typically screened for in the United States. WPW accounts for at least 1% of sudden death in athletes and occurs in at least 1 to 4.5 per 1000 adults and children, while symptoms tend to be asymptomatic [18].

As an athlete at St. John's University, I implore the efforts to further understand this heart condition and promote the screening for it at a young age. This research allows the opportunity to do just that by examining the characteristics of WPW by implementing it into our mode. Furthermore, the multi-class classification of other arrhythmias alike will be introduced in future iterations to correctly identify other common arrhythmias that are listed in the diagnostic classes in the dataset section.

4.2 Summary of Research

To inform a proper diagnosis and develop an understanding of Alzheimer's Disease progression, deep learning was used within a 3D convolutional architecture to provide an alternate approach for abnormality detection in MRI image data. The research provided offers attention to the efficacy and predictive power when detecting the progression of neurodegeneration within brain scans. The use of a volumetric data proves to yield classification accuracies comparable or better to other studies that ignore the third spatial dimension during training.

Alzheimer's disease continues to cause progressive issues within our society by overwhelming the physical and cognitive aspects of brain function that diminish a total sense of self. Although there is currently no cure, research in this field continues to grow and will hopefully facilitate future discoveries as well as the early detection of AD in MCI patients. Further research needs to be done to analyze the effectiveness of CNNs within medical imaging and neuroscience research, including the application between 3D architectures assisted by deep learning algorithms in its domain.

The research presented offers perspective to the use of deep learning and LSTM to identify arrhythmias in ECG recordings. Each row in the dataset represents a single ECG recording as input, allowing LSTM to identify patterns within the recording to denote abnormalities of multiple types. The binary classification of normal and abnormal arrhythmias was used to classify ECG recordings with an accuracy of 97%. The research found presents a discussion on the efficacy of using LSTM instead of CNNs to identify spectral features. Particularly in time series data, LSTM prove to be more effective in remembering informative features and forecast predictions in signals, compared to CNNs that classify images and no utilize memory within the network.

Heart conditions such as arrhythmias, heart disease, and strokes continue to present life-threatening challenges to individuals worldwide, promoting the attention to identification and treatment of such ailments. Arrhythmias can go unnoticed for years due to a lack of resources and the ability to recognize the prerequisite symptoms. This research offers perspective into the monitoring and detection of such arrhythmias due to their complexity and the ability to go undetected. Research in this domain can also be tested in the use of wearable monitors to detect arrhythmias in everyday activities, particularly in athletic-related activities.

4.3 Conclusion

In this study, we introduce a computational approach in an effort to solve real-world medical issues that continue to plague our society and are difficult to diagnose. The use of deep learning has offered an alternate form of detection and understanding of health conditions in a data-driven way, ushering in a new area of medical research that involves a perhaps more accurate method of diagnosis than leading practices. Technology has reached a point where the availability and amount of data have can finally be harvested to apply artificial learning techniques to benefit society, particularly in the healthcare domain. Research in this field continues to grow and will facilitate new discoveries to solve the issues previously mentioned and will promote a better understanding of health conditions alike.

REFERENCES

[1] Bousseljot, R.; Kreiseler, D.; and Schnabel, A. Nutzung der EKG Signal datenbank CARDIODAT der PTB über das Internet, 1995.

[2] Bondi, M. W., Edmonds, E. C., and Salmon, D. P. Alzheimer's Disease: Past, Present, and Future. *Journal of the International Neuropsychological Society* : JINS, 23(9-10), 818–831, 2017.

[3] Breiman, L. Bagging predictors. Machine learning, 24(2), 123-140, 1996.

[4] Carin, L. On artificial intelligence and deep learning within medical education. *Academic Medicine*, 95(11S), S10-S11, 2020.

[5] Christian, S.; Antonio, C.; Petronilla, B.; Gilardi, M. C.; Aldo, Q. and Isabella, C.; 2015. Magnetic resonance imaging biomarkers for the early diagnosis of Alzheimer's disease: a machine learning approach. Frontiers in neuroscience, 9, 307.

[6] Christoforou, C.; Constantinidou, F.; Shoshilou, P.; and Simos, P. 2013. Single-trial linear correlation analysis: application to characterization of stimulus modality effects. Frontiers in Computational Neuroscience, 7, 15.

[7] Christoforou, C.; Haralick, R.M.; Sajda, P.; and Parra, L. C. 2010. The bilinear brain: Towards subject-invariant analysis, *In 2010 4th International Symposium on Communications, Control and Signal Processing (ISCCSP*), pp. 1–6. IEEE, 2010.

[8] Christoforou, C.; Hatzipanayioti, A.; and Avraamides, M. 2018. Perspective-taking vs mental rotation: CSP-based single-trial analysis for cognitive process disambiguation. In Wang, S., Yamamoto, V., Jianzhong S., Yang Y., Jones, E., Iasemidis, L., Mitchell, T., (Eds.) *Proceedings of International Conference, Brain Informatics* (pp. 109-199). Arlington, TX, USA.

[9] Dupre, A.; Vincent, S.; and Iaizzo, P. A. Basic ECG Theory, Recordings, and Interpretation. Handbook of Cardiac Anatomy, Physiology, and Devices, 191–201.

[10] Elstein A. Clinical reasoning in medicine. In: Higgs J, ed. Clinical reasoning in the health professions. Oxford, England: Butterworth-Heinemann Ltd, 1995;49–59

[11] Gupta, H.; Jin, K. H.; Nguyen, H. Q.; McCann, M. T.; and Unser, M. CNN-Based Projected Gradient Descent for Consistent CT Image Reconstruction, in *IEEE Transactions on Medical Imaging, 37*(6), 1440-1453.

[12] Heart Disease and Stroke Statistics— 2021 Update: A Report From the American Heart Association.

[13] Kłosowski, G.; Rymarczyk, T.; Wójcik, D.; Skowron, S.; Cieplak, T.; and Adamkiewicz, P. The Use of Time-Frequency Moments as Inputs of LSTM Network for ECG Signal Classification, 2020.

[14] Kumar, A.; Kim, J.; Lyndon, D.; Fulham, M., and Feng, D.; 2017. An Ensemble of Fine-Tuned Convolutional Neural Networks for Medical Image Classification in *IEEE Journal of Biomedical and Health Informatics, 21*(1), 31-40.

[15] Kumar, J., Goomer, R., and Singh, A. K.; 2018. Long Short Term Memory Recurrent Neural Network (LSTM-RNN) Based Workload Forecasting Model For Cloud Datacenters.

[16] Lin, W.; Tong, T.; Gao, Q.; Guo, D.; Du, X.; and Yang, Y.; 2018. Alzheimer's Disease Neuroimaging Initiative. Convolutional neural networks-based MRI image analysis for the Alzheimer's disease prediction from mild cognitive impairment. *Frontiers in neuroscience*, 12, 777.

[17] Pan, D.; Zeng, A.; Jia, L.; Huang, Y.; Frizzell, T.; and Song, X.; 2020. Early Detection of Alzheimer's Disease Using Magnetic Resonance Imaging: A Novel Approach Combining Convolutional Neural Networks and Ensemble Learning. *Frontiers in neuroscience*, 14.

[18] Rao, A. L., Salerno, J. C., Asif, I. M., and Drezner, J. A.; 2014. Evaluation and management of wolff-Parkinson-white in athletes. *Sports health*, 6(4), 326–332.

[19] Razzak, M. I., Naz, S., and Zaib, A.; 2018. Deep learning for medical image processing: Overview, challenges and the future. *Classification in BioApps*, 323-350.

[20] Saadatnejad, S.; Oveisi, M.; and Hashemi, M.; 2019. LSTM-based ECG classification for continuous monitoring on personal wearable devices. *IEEE journal of biomedical and health informatics*, 24(2), 515-523.

[21] Sak, H.; Senior, A. W.; and Beaufays, F. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling.

[22] Saleh; Moussa; and John A Ambrose. "Understanding myocardial infarction." F1000Research vol. 7 F1000 Faculty Rev-1378. 3 Sep. 2018

[23] Salem, M.; Taheri, S.; and Yuan, J. 2018. "ECG Arrhythmia Classification Using Transfer Learning from 2- Dimensional Deep CNN Features," 2018 *IEEE Biomedical Circuits and Systems Conference* (BioCAS), pp. 1-4

[24] Singh, S.; Pandey, S. K.; Pawar, U.; & Janghel, R. R.; 2018. Classification of ECG Arrhythmia using Recurrent Neural Networks. Procedia Computer Science.

[25] Singh, S. P.; Wang, L.; Gupta, S.; Goli, H.; Padmanabhan, P.; and Gulyás, B.; 2020. 3D Deep Learning on Medical Images: A Review. Sensors, 20(18), 5097.

[26] Yang, H.; Xu, H.; Li, Q.; Jin, Y.; Jiang, W.; Wang, J.; and Wang, T.; 2019. Study of brain morphology change in Alzheimer's disease and amnestic mild cognitive impairment compared with normal controls. in *General Psychiatr*y, 32

[27] Zihlmann, M.; Perekrestenko, D.; & Tschannen, M. 2017. Convolutional recurrent neural networks for electrocardiogram classification. *In 2017 Computing in Cardiology* (CinC) (pp. 1-4). *IEEE*

Vita

Name	Ryan Hogan
Baccalaureate Degree	Bachelor of Science, St. John’s University, NYC, Major: Computer Science
Date Graduated	May 2020
Masters Degree	Master of Science, St. John’s University, NYC, Major: Computer Science with Machine Learning
Date Graduated	May 2021